

Macroscopic Dynamics in Large-Scale Data Networks

Jian Yuan^{1,2} and Kevin Mills¹

¹ Advanced Network Technologies Division, National Institute of Standards and Technology, Gaithersburg, Maryland 20899, USA

² Department of Electronic Engineering, Tsinghua University, Beijing 10084, China

1 Introduction

Modern society grows increasingly reliant on the Internet, a network of global reach that supports many services and clients. However, in such a large-scale distributed network, meeting quality-of-service requirements presents a difficult challenge because hotspots of network load move around and traffic anomalies arise unpredictably in space and time. In this chapter, we will demonstrate that observing network dynamics at a macroscopic level is likely to contribute to better network engineering and management.

The Internet is an enormous network of networks without central control or administration. Millions of computers around the world attach to the Internet through many autonomous regional networks of routers, which interconnect through backbone networks of routers in a distributed, hierarchical fashion. Internet computers exchange data with each other in units called packets, where each packet is accepted by a router, stored temporarily, and then forwarded on to a next router. This accept-store-and-forward cycle begins when a source computer transmits a packet to an entry router and continues until the packet is forwarded to its intended destination by an exit router. The Internet's design is guided by the end-to-end principle [1], which allocates simple functionality to routers, while pushing complexities of specific applications and of congestion avoidance mechanisms outside the network and into attached computers. The implication of this design principle is that Internet routers see no relationship among individual packets, while “end-to-end” protocols implemented in Internet-attached computers manage all state associated with data exchanges. The basic communication protocol of the Internet is called TCP/IP (Transmission-Control Protocol/Internet Protocol) [2]. IP is a “hop-by-hop” protocol used by source computers to inject packets into the Internet, and used by Internet routers to store-and-forward packets among multiple routers along a path, and then finally to forward the packet to its destination computer. TCP is an end-to-end protocol operating on logical connections between pairs of computers.

TCP includes a congestion-control algorithm to ensure that a sender does not transmit more data than the network can handle. The TCP congestion-control algorithm exhibits a self-organizing property: when a large number of logical

connections share the Internet, underlying interactions among the connections avoid router congestion simultaneously over varying spatial extent; however, the network-wide effects created by such interactions are difficult to determine. The spatial-temporal dynamics of Internet traffic is also difficult to characterize due to highly variable user demands and to unpredictable resource availability. Further, not every client using the Internet is honest or co-operative. For example, distributed denial-of-service (DDoS) attacks, which arise when large numbers of compromised computers send traffic simultaneously toward a victim (e.g., a web server or a router) [3], may intermittently disturb the normal operating condition of the Internet. All these sources of variability inhibit easy characterization of Internet-wide traffic dynamics.

We suspect that, where many globally distributed data flows simultaneously transit a large network, the self-organizing properties of the TCP congestion-control algorithm might lead to the emergence of collective behavior, as in other complex systems [4]. Collective emergent phenomena often can be identified when the behavior of an entire system appears more coherent and directed than the behavior of individual parts of the system. In this way, any single data flow across a large network would not face a totally random condition, but more likely would adapt itself to a steady collective state, in which the flow could make little change. If such an emergent collective property occurs in large networks, and if we can describe and visualize the associated patterns, then perhaps such knowledge can be used to improve global network performance and to increase resistance to subtly engineered DDoS attacks.

Since emergent coherent behavior exhibits a spatial-temporal dependence among collective data flows over a whole network, correlation might be key to describing emergent patterns. A number of empirical studies on traffic measurements have convincingly demonstrated that actual Internet traffic exhibits long-range dependence (LRD) [5-7], which implies the existence of nontrivial correlation structure at large timescales. However, in these studies, the LRD found in Internet traffic was not attributed to an emergent, spontaneous order at the macroscopic (whole network) level. Instead, these studies attributed LRD to the linear multiplexing of a large number of highly variable traffic sources [8]. This explanation apparently ignores any nonlinear relationships that might arise as collective flows compete for network resources (router buffers and link capacity) over space and time. To understand the potential collective effect in large-scale networks, we conducted our own studies to identify the reasons behind LRD traffic phenomena [9-11]. We found that network size has greater influence than other factors—e.g., high variability in traffic sources and choice of transport mechanism—on the *temporal* dynamics of network congestion. Our findings suggesting the importance of network size in generating emergent collective behavior led us to consider how we might examine both the spatial and temporal dynamics of network congestion.

Recently, graph wavelets have been proposed for *spatial* traffic analysis given knowledge of aggregate traffic measurements extracted at intervals over all links [12]. This method can provide a highly summarized view of traffic load throughout an entire network. There seems to be no stringent time limit for producing a snapshot of network-wide load with such spatial traffic analysis; however, spatial-temporal traffic analysis, which reveals the time-varying nature of spatial traffic, may have to perform in a timely manner. Currently, spatial-temporal traffic analysis presents practical difficulties, not only because large-scale distributed networks exhibit high-dimensional traffic data, but also because mining large amounts of data may strain

memory and computation resources in even the most advanced generation of desktop computers. Moreover, routers may be heavily utilized, and thus fail to collect and transfer data, often when the routers are of most interest (due to their congested nature). Given these practical constraints, it would be appealing to reduce the amount of data to transfer and process, while retaining the ability to observe spatial-temporal traffic dynamics. We believe that the emergence of collective behavior (with its associated global order) could be exploited to concisely capture spatial-temporal patterns with sparse observation points. In other words, if emergent behavior arises in a large network, then traffic will be correlated over wide space-time and, thus, might be characterized by sampling a small number of points. On the other hand, if network traffic exhibits little space-time correlation, then sampling a small number of points would not prove particularly revealing.

A recent study of correlations among data flows in a French scientific network, *Renater* [13], detected the signature of collective behavior. The *Renater* study uses methods from random matrix theory (RMT) to analyze cross-correlations between network flows. In essence, RMT compares a random correlation matrix—a correlation matrix constructed from mutually uncorrelated time series—against a correlation matrix for the data under investigation. Deviations between properties of the cross-correlation matrix from the investigation data and the correlations in the random data convey information about “genuine” correlations. In the case of the *Renater* study, the most remarkable deviations arise about the largest eigenvalue and its corresponding eigenvector. The largest eigenvalue is approximately a hundred times larger than the maximum eigenvalue predicted for uncorrelated time series. The largest eigenvalue appears to be associated with a strong correlation over the whole network. In addition, the eigenvector component distribution of the largest eigenvalue deviates significantly from the Gaussian distribution predicted by RMT. Further, the *Renater* study reveals that all components of the eigenvector corresponding to the largest eigenvalue are positive, which implies their collective contribution to the strong correlation. Since all network data flows contribute to the eigenvector, the eigenvector can be viewed as the signature of a collective behavior for which all flows are correlated.

In fact, the eigenvector corresponding to the largest eigenvalue provides an important clue, which led us to a novel method for observing spatial-temporal dynamics at the macroscopic level [14]. As the macroscopic pattern emerges from all adaptive behaviors of data flows in various directions, hotspots should be exposed through their correlation information, as the joining points of significantly correlated data flows. Note that the details of the components of the eigenvector of the largest eigenvalue reveal this information, with the larger components corresponding to the more correlated flows. Therefore, we define a weight vector by grouping eigenvector components corresponding to a destination routing domain together to build up information about the influence of the domain over the whole network. Contrasting weights of the weight vector against each other in space and time, we not only can summarize a network-wide view of traffic load, but can also locate hot spots, and can even observe how spatial traffic patterns change from one time period to the next.

Using this macroscopic-level method inevitably encounters issues of scale, that is, gathering data from numerous distributed measurement points, and consuming computation time and memory when analyzing data. The *Renater* study assumes complete information from all network connection points, which proves feasible because the *Renater* network contains only about 30 interconnected routers. We have

figured out how to scale down the coverage problem by exploiting an emergent collective phenomenon, called the correlation increase [14]. Correlation increases arise from *collective response* of the entire network to changes in traffic. This effect has already been observed in the framework of stock correlations, where cross-correlations become more pronounced during volatile periods as compared to calm periods [15]. Indeed, higher values of the largest eigenvalue occur during periods of high market volatility, which suggests strong collective behavior accompanies high volatility. This connection has value in our analysis because Internet traffic behavior appears to be nonstationary [16]. An increase in cross-correlation allows us to infer a shift in the spatial-temporal traffic pattern of large areas of interest outside those few areas where measurements are made. This approach could significantly reduce requirements for data, perhaps to the point where analysis could occur in real time.

In this chapter, we use simulation results to show how this innovation could succeed in a large TCP/IP network. We apply our technique to identify network hotspots and to expose large-scale DDoS attacks in our simulation environment. The rest of this chapter is structured as five sections. Section 2 delineates a simulation model we developed recently to study space-time characteristics of congestion in large networks, and to analyze system behavior as a coherent whole. In Section 3, we describe our technique for spatial-temporal traffic analysis. In Section 4, we show how our technique captures network-wide patterns shifting over time. Section 5 demonstrates the macroscopic effect of DDoS flooding attacks, and shows how our technique could provide significant information to detect and defend against such attacks. We present concluding remarks in Section 6.

2 Modeling a Large-scale TCP/IP Network

Network simulation plays a key role in building an understanding of network behavior. Choosing a proper level of abstraction for a model depends very much on the objective. Studying collective phenomena seems to require simulating networks with a large spatial extent. Appropriate models for such studies should also include substantial detail representing protocol mechanisms across several layers of functionality (e.g., application, transport, network, and link), yet must also be restricted in space and time in order to prove computationally tractable. Previously, we adopted a two-tier modeling approach that maintains the individual identity of packets, producing a full-duplex “ripple effect” at the packet level, and that also accommodates spatial correlations in a regular network structure [10, 11]. While our two-tier model has been applied successfully to qualitatively understand some traffic characteristics in large-scale networks [11, 14], some doubts exist about the realism inherent in the regular network structure of such a model. In this chapter, we retain the individual identity of packets but we replace the regular network structure of our previous two-tier model with a large-scale irregular topology chosen to resemble the topology of a real network.

2.1 Topology

Here, we transform our regular two-tier model into an irregular four-tier topology, as shown in Figure 1. (The host-computer tier is not shown in Figure 1.) While the

network core becomes heterogeneous and hierarchical, (tier-four) host-computer behavior remains homogeneous at the edge of the network. The (tier-one) backbone topology, including eleven (backbone) routers (A, B, ... K), resembles the original Abilene network, as described elsewhere [12]. Links between backbone routers have varying delays. For example, the longest link between backbone routers D and F has a 20-ms propagation delay; the shortest propagation delay (3 ms) exists on the link between backbone routers J and K. Forty (tier-two) subnet routers, represented by designators such as A1 and B2. Each subnet contains a variable number of (tier-three) leaf routers, such as A1a and B2b. Each leaf router supports an equal number (200 in this chapter) of (tier-four) source hosts, and a variable number (≤ 800 in this chapter) of (tier-four) receivers, activated on demand. Link capacities gradually increase from host (tier four) to backbone, with (tier-one) backbone links being hundreds of times faster than links to (tier-four) hosts.

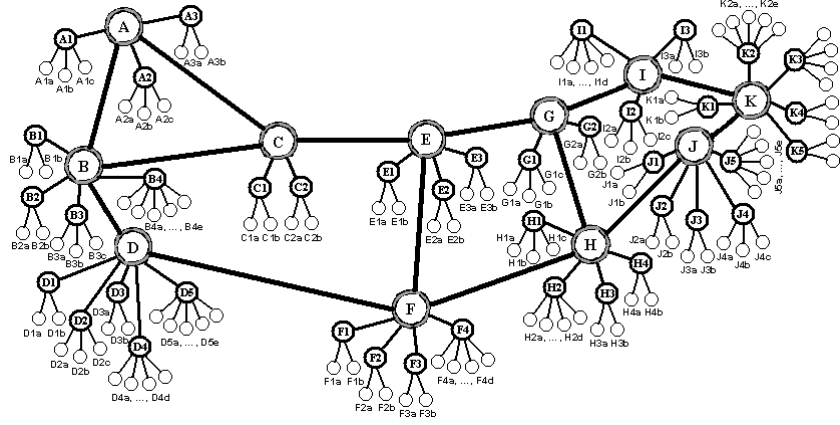


Figure 1: Four-tier simulation model with 11 (tier-one) backbone routers, 40 (tier-two) subnet routers, and 110 (tier-three) leaf routers. The 22,000 (tier-four) source hosts and the up to 88,000 (tier-four) receivers are not shown.

2.2 Traffic sources

There are a total of 22,000 sources in our model, which operates at the packet level. Each source models traffic generation as an ON/OFF process, which alternates between wake and sleep periods with average durations λ_{on} and λ_{off} , respectively. When awake, a source may send, subject to any restrictions imposed by TCP, one packet at each time-step to the source's attached leaf router. The packet will be placed at the end of the router's queue. At the beginning of each ON period, a destination receiver is chosen randomly from among leaf routers other than the local leaf routers, i.e., all flows must transit through at least one backbone link. When sleeping, the source does not generate new packets at each time-step. ON/OFF sources provide a convenient model of user behavior.

Empirical measurements on the Internet observe a heavy-tailed distribution of transferred file sizes [7]. Here, we use the Pareto distribution for both ON and OFF processes with the same shape parameter α [11]. In this chapter, $\lambda_{on} = 50$, $\lambda_{off} = 5000$ and $\alpha = 1.5$.

2.3 Routers

Packets, the basic unit of transmission on TCP/IP networks, contain destination addresses used by routers to correctly forward and source addresses used by receivers to identify the destination address for reply packets. To store and forward packets, which in our model travel a constant, shortest path between a source-destination pair for each flow, all routers maintain a queue of limited length (160 packets/router here), where arriving packets are stored until they can be processed: first-in, first-out. For convenience, in this chapter we assume that every discrete simulation time-step is 1 millisecond. However, each leaf router, subnet router, or backbone router can in turn forward 5, 20, or 160 packets during one millisecond. This simulates capacity differences among various link classes from leaf-access to backbone in a hierarchically structured network. With such parameter settings, simulated backbone links are very lightly loaded.

We select several subnet routers as observation points, e.g., B4, D5, F4, I1, and J5, which record all outbound flows to destination leaf routers. In this chapter, we assume that a central collector reliably receives a continuous stream of measured data from observation points in time to perform analysis for our various experiments.

3 Representing Macroscopic-level Traffic Dynamics

In this section, we discuss briefly our approach to represent traffic dynamics at a macroscopic-level. First, we describe how we represent network flow data. Second, we outline our use of cross-correlation analysis. Finally, we depict our technique to summarize network-wide traffic load using a weight vector.

3.1 Representing network flow data

Assume that there are N leaf routers, interconnecting through subnet routers and backbone routers to form a large-scale distributed network, where L subnet routers are deployed as observation points to log outbound traffic. First, we need to represent packets flowing between distinct source-destination pairs at each sampling interval. Let $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)^T$ denote the flow vector of corresponding packet counts, observed in L subnets during a given time interval. Each element of this flow vector is itself a vector defining the number of packets flowing into the corresponding leaf router from each of the observation subnets in the distributed network. The method to obtain all flow variables in this vector is to first enumerate all the destination leaf routers and then the observation posts by 1 to L , and group these indices by leaf router: the subnets sending to the first leaf router in the first block, \mathbf{x}_1 , and those sending to the second leaf router in the second block, \mathbf{x}_2 , and so forth. Thus, we form \mathbf{x} with subvectors in the order $\mathbf{x}_1 = (x_{11}, x_{21}, \dots, x_{L1})^T$, $\mathbf{x}_2 = (x_{12}, x_{22}, \dots, x_{L2})^T$, \dots , $\mathbf{x}_N = (x_{1N}, x_{2N}, \dots, x_{LN})^T$, where x_{ij} represents packet flow from the i th observation point (i

$=1, 2, \dots, L$) to the j th leaf router ($j=1, 2, \dots, N$). Each flow variable x_{ij} is normalized as f_{ij} by its mean m_{ij} and standard deviation σ_{ij} ,

$$f_{ij} = (x_{ij} - m_{ij}) / \sigma_{ij}. \quad (1)$$

Then, the normalized flow vector \mathbf{f} , corresponding to \mathbf{x} , comprises N normalized subvectors, \mathbf{f}_k ($k=1, 2, \dots, N$), where each subvector is formed from normalized flow variables f_{ik} ($i \leq L$ and $k \leq N$).

3.2 Cross-correlation analysis

Cross-correlation analysis is a tool commonly used to analyze multiple time series. We can compute the equal-time cross-correlation matrix \mathbf{C} with elements

$$C_{(ij)(kl)} = \langle f_{ij}(t) f_{kl}(t) \rangle, \quad (2)$$

which measures the correlation between f_{ij} and f_{kl} , where $\langle \dots \rangle$ denotes a time average over the period studied. The cross-correlation matrix is real and symmetric, with each element falling between -1 and 1 . Positive values indicate positive correlation, while negative values indicate an inverse correlation. A zero value denotes lack of correlation.

We can further analyze the correlation matrix \mathbf{C} through eigenanalysis [17]. The equation

$$\mathbf{C}\mathbf{w} = \lambda\mathbf{w} \quad (3)$$

defines eigenvalues and eigenvectors, where λ is a scalar, called the eigenvalue. If \mathbf{C} is a square K -by- K matrix, e.g., $K = L(N-1)$ here, then \mathbf{w} is the eigenvector, a nonzero K by 1 vector (a column vector). Eigenvalues and eigenvectors always come in pairs that correspond to each other. This eigenvalue problem has K real eigenvalues, some of which may repeat. An eigenvector is a special kind of vector for the matrix it is associated with, because the action of the underlying operator represented by the matrix takes a particularly simple form on the eigenvector input: namely, simple rescaling by a real number multiple. The eigenvector \mathbf{w}^I corresponding to the largest eigenvalue λ_I often has special significance for many applications. There are various algorithms for the computation of eigenvalues and eigenvectors [17]. Here, we exploit the MATLAB `eig` command, which uses the QR algorithm to obtain solutions [18].

3.3 Defining the weight vector

The cross-correlation matrix contains within itself information about underlying interactions among various flows. The components of the eigenvector \mathbf{w}^I of the largest eigenvalue λ_I represent the corresponding flows' influences on macroscopic behavior, abstracted from the matrix \mathbf{C} into the pair $(\lambda_I, \mathbf{w}^I)$. The eigenvector \mathbf{w}^I comprises N subvectors, i.e., $\mathbf{w}^I = (\mathbf{w}_1^I, \mathbf{w}_2^I, \dots, \mathbf{w}_N^I)^T$. The k th subvector \mathbf{w}_k^I , corresponding to the k th leaf router, is formed from components w_{ik}^I ($i \leq L$ and $k \leq N$) representing the i th observation point's contribution to the k th leaf router. We consider the square of each component, $(w_{ik}^I)^2$, instead of w_{ik}^I itself because $\sum_{i,k} (w_{ik}^I)^2 = 1$ [19]. We define the weight S_k ($k = 1, 2, \dots, N$) to be the sum of all $(w_{ik}^I)^2$ in the k th subvector \mathbf{w}_k^I , i.e.,

$$S_k = \sum_i^L (w_{ik}^1)^2. \quad (4)$$

S_k represents the relative strength of the contributions of the flows towards the k th leaf router. Thus, the knowledge of weight vector $\mathbf{S} = (S_1, S_2, \dots, S_N)$ across varying k constitutes one summary view of network-wide traffic load. The evolving pattern of spatial-temporal correlation might allow us to infer where and when network congestion emerges.

4 Capturing Shifting Spatial-temporal Patterns

Internet access is never evenly distributed. Flash-crowds are quite common. Hot spots might develop and break up more quickly than the network could be re-provisioned to respond. However, capturing the movement of hot spots seems very difficult. Here, we try to use our technique to observe the macroscopic dynamics of such phenomena.

To deliberately induce congestion, we let one selected leaf router have an additional five percent probability for selection as the destination domain. This is a natural way to change the network-wide traffic demand at longer timescale. Figure 2 depicts a change in congestion in leaf routers. The vertical axis represents the congested location within 11 backbone-router zones, each of which denotes the subset of leaf routers therein. At first, leaf router H4b is congested (up until time, t , is 400 s). From $t = 400$ s, C2b is selected as a new location to induce congestion. This congestion-induction technique offers an easily interpreted framework to analyze spatial-temporal pattern shifts driven by varying traffic demand.




Figure 2: Congested location changing over time from leaf router H4b to leaf router C2b.

4.1 Timescale of interest

When focusing on network-wide behavior, the timescale of interest should not be fine-grained. The microscopic fluctuations observed at shorter timescales usually reflect local details, while the driving force of traffic demand seems to vary over much longer timescales. The timescale of interest in our experiments appears at a middle range, similar to the concept of a critical timescale beyond which traffic fluctuation is supposed to exhibit greater influence than microscopic fluctuations [20]. At this middle timescale, macroscopic (coherent) behavior emerges as a connecting link between short-range microscopic fluctuations and the longer-range driving force of variations in traffic demand. This coherence is expected to emerge as a result of adaptive behaviors among data flows in different directions, but then to continue to shift its spatial-temporal pattern under the force of traffic demand.

We first form an observation system of eleven points ($L = 11$) by selecting one subnet router in each backbone-router zone, instead of observing all subnets in all backbone-router zones. We observe, at a granularity of 200ms, every fine-grain flow from these subnet routers to every leaf router ($N = 110$). (In a subsequent section, we will try to further reduce observation points.)

Now, we calculate the weight vector S with M data points ($M = 200$ in this chapter), which span a first period ($M/2$ points) and a second period ($M/2$ points). Selecting an appropriate data length for analysis might be largely considered a trial-and-error process (or the subject of future work). Here, we selected $M = 200$, which seemed to work fine. We tried $M = 100$ and 300 , which confirmed a data length of 200 more suitable for our experiments. Two weight vectors are calculated at the aggregated levels $T = 0.4$ s and $T = 2$ s, and shown respectively in Figure 3(a) and 3(b). The weight vector with $T = 2$ s shows two prominent weights at leaf routers C2b and H4b (S_{C2b} and S_{H4b}), revealing the network-wide pattern of congestion arising in these two domains. However, the pattern does not appear when $T = 0.4$ s. To clarify the role of timescale here, we further show the sum of S_{C2b} and S_{H4b} at different aggregated levels in Figure 3(c). We find that the sum of S_{C2b} and S_{H4b} gradually increases as T increases, up until about $T = 2$ s.

4.2 Increased correlation

Figure 4(a) shows the sum of S_{C2b} and S_{H4b} , which is calculated with $T = 1.6$ s and with the time window, MT ($= 200 \times 1.6$ s $= 320$ s), sliding ahead every 16 s. The corresponding λ_1 shows in Figure 4(b). The time axis indicates the end of the moving time window. The sum of S_{C2b} and S_{H4b} , and the largest eigenvalue λ_1 undulate almost in the same way, reaching higher values during the period of pattern shifting than during calm periods. The increased correlation in the simulation data emerges gradually after the second period starts, spreading the varying traffic demand to the entire network. During this transient period, flows in different directions have to adapt their behaviors to the changing congestion, and the flows continue to react to each other until they reach collectively a new coherent pattern.

With the measurement and analysis method, as outlined above in Section 3, applied at the appropriate timescale, as cross-correlations become more pronounced, traffic patterns over the whole system become more visible. In the remaining experiments, described below, we focus on macroscopic dynamics at the timescale $T = 2$ s.

4.3 Spatial-temporal pattern

It might prove feasible to design sample-based techniques suitable to identify network-wide patterns that remain invariant for a long time. However, when traffic demands vary over a large dynamic space-time range, these same techniques could fail to detect the more quickly changing patterns. By taking advantage of increased correlation arising over volatile periods, we might be able to use a sample-based version of our proposed method to identify shifting network-wide congestion patterns. In the following, we use only measurements from five (i.e., $L = 5$) subnet routers (B4, D5, F4, I1, and J5) to perform our analysis.

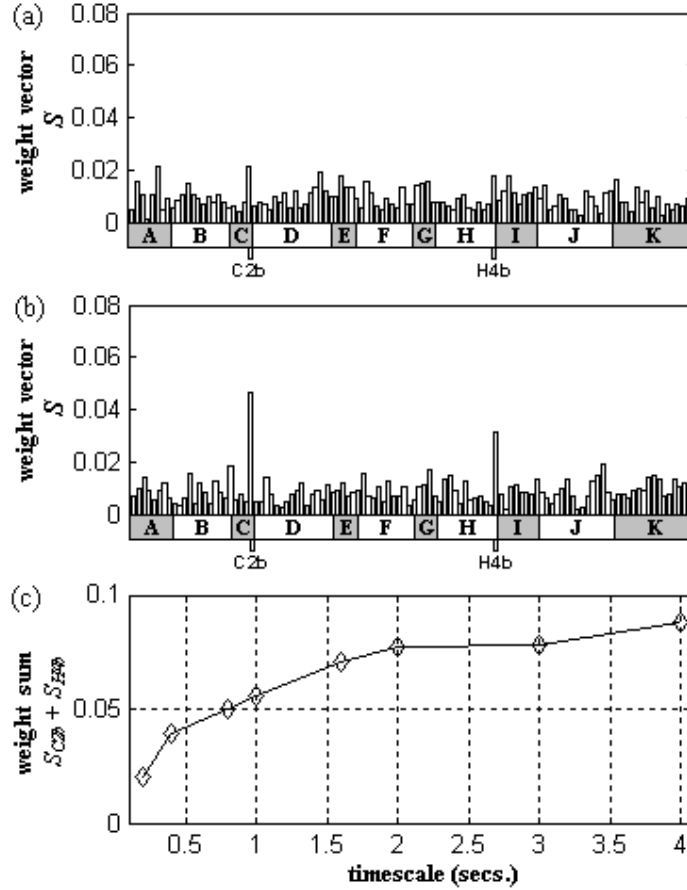


Figure 3: Two weight vectors at $T = 0.4$ s (a) and $T = 2$ s (b), and (c) the sum of S_{C2b} and S_{H4b} changing at different timescales with $L = 11$

To show how the spatial traffic pattern changes, we calculate the weight vector S using M data points within a moving time window MT from one time period to the next. Figure 5 shows the weight vector S evolving with $T = 2$ s and with the time

window MT ($= 200 \times 2 \text{ s} = 400 \text{ s}$) sliding ahead every 10 s. The time axis indicates the end of the moving time window. We can see the enhanced weights of C2b and H4b in the shifting spatial-temporal pattern. While the new congestion appears at C2b, the existing congestion at H4b, which was indistinguishable during the previous calm period, also exposes itself to the weight vector. The five observation points, which are not near C2b or H4b, really “sense” by themselves the gentle load fluctuation of these two leaf routers. The load wave seems to bring about a collective response in the entire network. This indicates that network-wide traffic appears correlated, and that spatial-temporal dynamics evolves as a coherent whole at some appropriate timescale. Therefore, macroscopic-level observation appears to provide significant information that could be exploited to achieve better network engineering and management.




Figure 4: (a) the sum of S_{C2b} and S_{H4b} , and (b) the largest eigenvalue λ_l with $T = 1.6 \text{ s}$ and $L = 11$

With fewer observation points, the increased correlation during transient periods is very helpful for capturing the network-wide (spatial) pattern of traffic shifting over time. While both S_{C2b} and S_{H4b} become enhanced during periods of shifting pattern, we know, as shown in Figure 2, that the congestion on C2b will persist, and that H4b will gradually recover its normal condition. If we need distinguish among routers with increasing and diminishing congestion, then other techniques, such as active probing for bandwidth or delay, might be applied to specific targets identified by our passive method of network-wide observation.

In larger networks, such as the Internet, it is very difficult, if not impossible, to observe the spatial-temporal pattern of congestion over the whole top tier, which

encompasses on the order of 10,000 autonomous systems. As discussed in the next section, spatial aggregation, e.g., from the leaf-router to subnet-router level, can help to implement a coarser space and time observation. First, however, we try to observe only a selected subset of the top network tier for the case of shifting congestion illustrated in Figure 2, while still including leaf-router details. We use measurements from the same five routers (subnets B4, D5, F4, I1, and J5) as before to form a spatial-temporal pattern over only three backbone-router zones of C, D, and E (comprising 26 leaf routers). Figure 6 shows the spatial-temporal pattern of the three regions, and reveals the congestion arising in C2b. This result suggests that our technique might provide a useful means to observe spatial-temporal dynamics in selected networks in a timely manner.




Figure 5: The spatial-temporal pattern evolving with $T = 2$ s and $L = 5$




Figure 6: The spatial-temporal pattern observed in three backbone-router zones, C, D and E, with $L = 5$

5 Monitoring DDoS Flooding Attacks

Distributed denial of service (DDoS) attacks present a very serious threat to the stability of the Internet. By simply exploiting the tremendous asymmetry between the large-scale distributed network resources and local capacities at the victim, a flooding-based DDoS attack can build up an intended congestion very quickly near an attacked target. DDoS attacks use forged source addresses and other techniques [21] to conceal the locations of the true attack sources; thus DDoS attacks are among the most difficult to detect and stop. Today's Internet infrastructure is extremely vulnerable to such large-scale coordinated attacks, which may easily and effectively remove an attack victim from the Internet, even without exploiting any particular vulnerabilities in network protocols or weaknesses in system design, implementation, or configuration.

To avoid congestion in the Internet, all flows under end-to-end controls adapt themselves in a self-organized, distributed manner. This adaptive behavior of flows in different directions plays a crucial role in keeping the Internet stable and in forming macroscopic traffic patterns. During a DDoS attack, the attack sources do not honor the normal end-to-end congestion control algorithms; rather, they overwhelm the intended victim, causing legitimate, well-behaved flows to back off, and then ultimately to starve. In addition, large-scale DDoS attacks also impair transit traffic flows, which happen to share a portion of the congested network. Such network-wide phenomena might show themselves in shifting patterns of spatial-temporal traffic.

5.1 Modeling DDoS attacks

To observe the macroscopic effect of DDoS attacks, we arrange 50 attack sources in our simulation model, which are distributed uniformly throughout the network. We enable our attack sources to launch constant-rate attacks collectively or using a subgroup technique (described further below). In our experiments, there are a total of 22,000 source nodes, and more than 10,000 simultaneously active TCP connections; thus, DDoS flows cannot be easily identified from the legitimate background traffic.

Usually, DDoS attacks directed against the network infrastructure can lead to more widespread damage than those directed against individual web servers. Here, one leaf router (I1a) will be the attack target. Routers under attack may fail to collect and transfer measurement data. Usually, it is difficult to monitor areas of interest without obtaining measurements from those areas. However, our analysis technique provides the ability to monitor areas of interest without such local measurements. We assume in our experiments that the attack on I1a disables the observation point deployed at the subnet-router I1; thus, we perform our analysis using data from only four observation points (B4, D5, F4 and J5; $L = 4$).

5.2 Constant rate attack

Constant rate, the simplest attack technique, is typical of known DDoS attacks. We arrange for all the 50 attack sources to launch constant-rate attacks collectively (that is, simultaneously). Here, we do not have the attack sources generate attack packets with full force [22], so that they cannot be easily identified through attack intensity at the source or in intermediate networks. We assume that the variable H represents the

intensity of an attack source. Since sources can only create one packet every millisecond, the maximum attack rate is one packet per millisecond, i.e., $H \leq 1$ (packet/ms). We experiment with a constant-rate DDoS attack where $H = 1/10$, that is, each attack source creates one attack packet for every 10 milliseconds beginning from $t_0 = 500$ s.

Figure 7 shows the weight vector S evolving with $T = 2$ s and with the time window $MT (= 200 \times 2 \text{ s} = 400 \text{ s})$ sliding ahead every 10 s. We find that the attack really leads to a network-wide shift of spatial-temporal correlation, and the congestion on the victim (I1a) reveals itself at the enhanced weight of I1a. Since we observe this phenomenon and get the time and location of the attack without any help from the suffering victim, the network-wide monitoring could be used to activate specific detection and filtering mechanisms to isolate and stop the attack flows.




Figure 7: The spatial-temporal pattern of the constant-rate attack with $H = 1/10$ and $L = 4$

We also can observe the spatial-temporal pattern of the constant-rate attack at the subnet level by spatially aggregating the destinations of network flow at the subnet level from current measurements at the leaf-router level. Figure 8 shows such a coarser observation, where the weight vector S evolves with $T = 2$ s and with the time window $MT (= 200 \times 2 \text{ s} = 400 \text{ s})$ sliding ahead every 20 s. Here, we can find that the constant-rate DDoS attack against I1a also results in the congestion on the subnet I1. With a lower computing time requirement, the coarser observation at this upper level still reveals a very useful picture of spatial-temporal dynamics.

5.3 Subgroup attack

Attackers constantly modify attack dynamics to evade detection. Attack dynamics can be made very sophisticated should an attacker desire. For example, next we divide the 50 attack sources into three subgroups, which are distributed separately in the left, the middle and the right parts of the larger network. Once the attack starts at $t_0 = 500$ s, one of the three subgroups is always active so that the victim experiences continuous denial of service [21]. Given the dynamic nature of such a coordinated attack, it is extremely hard to detect where attack packets originate, and to stop them at intermediate or source networks to reduce overall congestion and increase resources available to legitimate traffic.




Figure 8: The spatial-temporal pattern of $H = 1/10$ constant-rate attack at the subnet level with $L = 4$

Figure 9 shows the weight vector S evolving with $T = 2$ s and with the time window $MT (= 200 \times 2 \text{ s} = 400 \text{ s})$ sliding ahead every 10 s. We find that the subgroup attack reveals itself in the shifting spatial-temporal pattern. Comparing Figures 7 and 9, we find that for our analysis technique the dynamic nature of the subgroup attack seems advantageous, because the increased correlation induced by shifts in attack traffic keeps the weight of the victim I1a salient over a longer time range.

During the subgroup attack, we also observed a smaller portion of the larger distributed network, aggregated at the subnet level. Figure 10 shows the spatial-temporal pattern of five backbone-router zones (from G to K), where the weight vector S evolves with $T = 2$ s and with the time window $MT (= 200 \times 2 \text{ s} = 400 \text{ s})$ sliding ahead every 20 s, revealing the congestion arising in the subnet I1. The effects of the subgroup attack remain evident, while the aggregated, subnet-level observation

of only a portion of the network requires less computing time than for the case of Figure 8.




Figure 9: The spatial-temporal pattern of the subgroup attack with $H = 1/5$ and $L = 4$




Figure 10: The spatial-temporal pattern of the subgroup attack, observed in five zones of backbone routers from G to K with $L = 4$

6 Concluding remarks

In large-scale networks, such as the Internet, spatial-temporal correlations emerge from interactions among adaptive transport connections and from variations in user demands. By exploring the collective dynamics of large-scale networks, we seek ways to understand spatial-temporal correlations. We realize that capturing macroscopic patterns in correlations over time may help us to understand shifting traffic patterns, to identify operating conditions, and to reveal traffic anomalies.

Analyzing spatial-temporal characteristics of traffic in large-scale networks requires both a suitable analysis method and a means to reduce the amount of data that must be collected. In particular, routers may be heavily utilized or under DDoS attack, and thus fail to collect and transfer data, but often also happen to be the parts of interest to monitor (due to their congested nature). In this chapter, we describe a novel technique that provides a useful way to observe network-wide congestion patterns shifting over time. To illustrate this technique and its potential promise, we reported results from some simulation experiments.

We applied this technique successfully to identify network hotspots induced deliberately in a large-scale network. In particular, the effect of transient periods helped us to capture the network-wide traffic pattern shifting over time. We indicated that the spatial-temporal dynamics of network traffic appears as a coherent whole at an appropriate timescale.

We demonstrated how to use this novel technique to expose large-scale distributed denial-of-service (DDoS) attacks. We find that DDoS flooding attacks lead to a network-wide shift in spatial-temporal correlation, and that congestion on the attack victim reveals itself in these spatial-temporal patterns. The macroscopic effect of DDoS attacks can provide significant information about where and when a DDoS attack might be underway, and could trigger further detection and filtering without any information from the attack victim. In particular, we find that the dynamic nature of the (more stealthy) subgroup attack seems to be an advantage in revealing the victim's plight, because increased variation in traffic patterns lead to increased correlation, which is exploited by our analysis technique.

Since observing the whole Internet in detail is impractical, we suggested a means to efficiently observe selective portions in detail, or to apply spatial aggregation to observe larger-scale networks with less detail. In either case, our analysis method lowers computing time requirements, while revealing shifting traffic patterns over both space and time. If proven successful when applied to real network measurement data, our proposed technique could become a powerful tool to monitor spatial-temporal behavior network-wide in real time, and could ultimately contribute to improvements in network engineering and management.

References

- [1] J. Saltzer, D. Reed, and D. Clark, End-to-end arguments in system design, *ACM Trans. Computer System*, 2(4), pp. 277-288, November 1984.
- [2] W. R. Stevens, *TCP/IP Illustrated*, Vol. 1, Addison-Wesley Pub. Co., Reading, MA, 1994.

- [3] D. Moore, G. Voelker, and S. Savage, Inferring Internet denial of service activity, In Proceedings of the USENIX Security Symposium, Washington, DC, USA, August 2001.
- [4] Yaneer Bar-Yam, The Dynamics of Complex Systems (Studies in Nonlinearity), ISBN 0-201-55748-7, August 1997.
- [5] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, On the self-similar nature of Ethernet traffic, in: Proc. ACM SIGCOMM '93, 183-193, 1993.
- [6] V. Paxson and S. Floyd, Wide-area traffic: The failure of Poisson modeling, in: Proc. ACM SIGCOMM '94, 257-268, 1994.
- [7] M. E. Crovella and A. Bestavros, Self-similarity in world wide web traffic: Evidence and possible causes, in: Proceedings of the 1996 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, May 1996.
- [8] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level, in: Proc. ACM SIGCOMM '95, 100-113, 1995.
- [9] J. Yuan, Y. Ren, and X. Shan, Self-organized Criticality in a Computer Network Model, Physical Review E, 61(2), 1067-1071, 2000.
- [10] J. Yuan, K. Mills, Exploring Collective Dynamics in Communication Networks, Journal of Research of the National Institute of Standards and Technology, 107 (2), 179-191, 2002.
- [11] J. Yuan and K. Mills, Implication of Internet Traffic Characteristics for Network-Adaptive Distributed Systems, submitted, 2003.
- [12] M. Crovella and E. Kolaczyk, Graph Wavelets for Spatial Traffic Analysis, in: Proceedings of IEEE Infocom 2003, San Francisco, CA, USA, April 2003.
- [13] M. Barthelemy, B. Gondran, and E. Guichard, Large scale cross-correlations in Internet traffic, Physical Review E 66 (2002) 056110.
- [14] J. Yuan, K. Mills, A cross-correlation based method for spatial-temporal traffic analysis, submitted, 2003.
- [15] V. Plerou, *et al.*, Random matrix approach to cross correlations in financial data, Physical Review E 65 (2002) 066126.
- [16] K. Thompson, G. J. Miller, and R. Wilder, Wide-Area Internet Traffic Patterns and Characteristics, IEEE Network 11 (6) (1997) 10-23.
- [17] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2000.
- [18] The MathWorks, Inc., Natick, MA, USA, MATLAB User's Guide, 1998.
- [19] K. I. Goh, B. Kahng, and D. Kim, Spectra and eigenvectors of scale-free networks, Physical Review E 64 (2001) 051903.
- [20] M. Grossglauser and D. Tse, A time-scale decomposition approach to measurement-based admission control, In: Proceedings of IEEE Infocom '99, pp. 1539-1547, New York, NY, March 1999.
- [21] J. Yuan, K. Mills, Monitoring the macroscopic effect of DDoS flooding attacks, in preparation, 2003.
- [22] J. Mirkovic, G. Prier and P. Reiher, Attacking DDoS at the Source, Proceedings of ICNP 2002, pp. 312-321, Paris, France, November 2002.